

PATENT APPLICATION

**METHODS AND APPARATUS FOR IMPLEMENTING HOME
AGENT REDUNDANCY**

Inventors:

Kent K. Leung
2447 Villa Nueva Way
Mountain View, CA 94040
Citizenship: United States

Madhavi W. Chandra
113 Holmhurst Court
Apex, NC 27502
Citizenship: United States

Assignee:

Cisco Technology, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
A corporation of California

Status: Large Entity

Prepared by:

BEYER, WEAVER & THOMAS, LLP

METHODS AND APPARATUS FOR IMPLEMENTING HOME AGENT REDUNDANCY

5

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to Mobile IP network technology. More particularly, the present invention relates to methods and apparatus for enabling Home Agent redundancy.

2. Description of the Related Art

Mobile IP is a protocol which allows laptop computers or other mobile computer units (referred to as "Mobile Nodes" herein) to roam between various sub-networks at various locations -- while maintaining internet and/or WAN connectivity. Without Mobile IP or related protocol, a Mobile Node would be unable to stay connected while roaming through various sub-networks. This is because the IP address required for any node to communicate over the internet is location specific. Each IP address has a field that specifies the particular sub-network on which the node resides. If a user desires to take a computer which is normally attached to one node and roam with it so that it passes through different sub-networks, it cannot use its home base IP address. As a result, a business person traveling across the country cannot merely roam with his or her computer across geographically disparate network

segments or wireless nodes while remaining connected over the internet. This is not an acceptable state-of-affairs in the age of portable computational devices.

To address this problem, the Mobile IP protocol has been developed and implemented. An implementation of Mobile IP is described in RFC 2002 of the IP Routing for Wireless/Mobile Hosts Working Group, C. Perkins, Ed., October 1996. Mobile IP is also described in the text "Mobile IP Unplugged" by J. Solomon, Prentice Hall. Both of these references are incorporated herein by reference in their entireties and for all purposes.

The Mobile IP process and environment are illustrated in FIG. 1A. As shown there, a Mobile IP environment 2 includes the internet (or a WAN) 4 over which a Mobile Node 6 can communicate remotely via mediation by a Home Agent 8 and a Foreign Agent 10. Typically, the Home Agent and Foreign Agent are routers or other network connection devices performing appropriate Mobile IP functions as implemented by software, hardware, and/or firmware. A particular Mobile Node (e.g., a laptop computer) plugged into its home network segment connects with the internet through its designated Home Agent. When the Mobile Node roams, it communicates via the internet through an available Foreign Agent. Presumably, there are many Foreign Agents available at geographically disparate locations to allow wide spread internet connection via the Mobile IP protocol. Note that it is also possible for the Mobile Node to register directly with its Home Agent.

As shown in FIG. 1A, Mobile Node 6 normally resides on (or is "based at") a network segment 12 which allows its network entities to communicate over the internet 4 through Home Agent 8 (an appropriately configured router denoted R2). Note that Home Agent 8 need not directly connect to the internet. For example, as

shown in FIG. 1A, it may be connected through another router (a router R1 in this case). Router R1 may, in turn, connect one or more other routers (e.g., a router R3) with the internet.

Now, suppose that Mobile Node 6 is removed from its home base network segment 12 and roams to a remote network segment 14. Network segment 14 may include various other nodes such as a PC 16. The nodes on network segment 14 communicate with the internet through a router which doubles as Foreign Agent 10. Mobile Node 6 may identify Foreign Agent 10 through various agent solicitations and agent advertisements which form part of the Mobile IP protocol. When Mobile Node 6 engages with network segment 14, it composes a registration request for the Home Agent 8 to bind the Mobile Node's current location with its home location. Foreign Agent 10 then relays the registration request to Home Agent 8 (as indicated by the dotted line "Registration"). During the registration process, the Home Agent and the Mobile Node 6 may then negotiate the conditions of the Mobile Node's attachment to Foreign Agent 10. For example, the Mobile Node 6 may request a registration lifetime of 5 hours, but the Home Agent 8 may grant only a 3 hour period. Therefore, the attachment may be limited to a period of time. When the negotiation is successfully completed, Home Agent 8 updates an internal "mobility binding table" which links the Mobile Node's current location via its care-of address (e.g., a collocated care-of address or the Foreign Agent's IP address) to the identity (e.g., home address) of Mobile Node 6. Further, if the Mobile Node 6 registered via a Foreign Agent, the Foreign Agent 10 updates an internal "visitor table" which specifies the Mobile Node address, Home Agent address, etc. In effect, the Mobile Node's home base IP address (associated with segment 12) has been binded to the care-of address such as the Foreign Agent's IP address (associated with segment 14).

Now, suppose that Mobile Node 6 wishes to send a message to a Corresponding Node 18 from its new location. An output message from the Mobile Node is then packetized and forwarded through Foreign Agent 10 over the internet 4 to Corresponding Node 18 (as indicated by the dotted line "packet from MN")

5 according to a standard internet protocol. If Corresponding Node 18 wishes to send a message to Mobile Node -- whether in reply to a message from the Mobile Node or for any other reason -- it addresses that message to the IP address of Mobile Node 6 on sub-network 12. The packets of that message are then forwarded over the internet 4 and to router R1 and ultimately to Home Agent 8 as indicated by the dotted line 10 ("packet to MN(1)"). From its mobility binding table, Home Agent 8 recognizes that Mobile Node 6 is no longer attached to network segment 12. It then encapsulates the packets from Corresponding Node 18 (which are addressed to Mobile Node 6 on network segment 12) according to a Mobile IP protocol and forwards these encapsulated packets to a "care of" address for Mobile Node 6 as shown by the dotted 15 line ("packet to MN(2)"). The care-of-address may be, for example, the IP address of Foreign Agent 10. Foreign Agent 10 then strips the encapsulation and forwards the message to Mobile Node 6 on sub-network 14. The packet forwarding mechanism implemented by the Home and Foreign Agents is often referred to as "tunneling."

FIG. 1B illustrates a significant problem with the Mobile IP system 2. If Home 20 Agent 8 fails or otherwise become inoperative (due to a power failure, rebooting, scheduled maintenance, etc.), Mobile Node 6 is left without the ability to (1) receive new internet messages addressed to it at network segment 12 and (2) register with other Foreign Agents. In effect, Mobile Node 6 is cut-off from internet connection when Home Agent 8 goes down. This problem may extend to other Mobile Nodes supported by

Home Agent 8. Often, a given Home Agent will be responsible for servicing numerous Mobile Nodes which may be based at sub-network 12, and hence, all registered Mobile Nodes will lose connectivity.

As shown in FIG. 1B, when Home Agent 8 fails, not only is network segment 12 disconnected from the internet, but Mobile Nodes at remote locations are also blocked from the registration and packet receipt functions of Mobile IP. In some networks, there may be other routers connecting segment 12 to the internet. Such additional routers would allow fixed hosts on the segment to maintain their internet connections but would not allow remote access to Mobile Nodes. Similarly, if Foreign Agent 10 should fail, all Mobile Nodes visiting sub-network 14 lose connections, even though there may be other routers on that sub-network.

A redundancy protocol known as Hot Standby Router Protocol ("HSRP") is widely used to back up primary routers for a network segment. In HSRP, a "standby" or "secondary" router is designated as the back-up to an "active" or "primary" router. The standby router is linked to the network segment or segments serviced by the active router. The active and standby routers share a "virtual IP address" and possibly a "virtual Media Access Control (MAC) address." All internet communication to and from the relevant sub-network employs the virtual IP and MAC addresses. At any given time, the active router is the only router adopting the virtual addresses. Then, if the active router should cease operation for any reason, the standby router immediately takes over its load (by adopting the virtual addresses). Further details of HSRP can be found in RFC 2281, "Cisco Hot Standby Router Protocol (HSRP)" by T. Li, B. Cole, P. Morton, and D. Li and in US Patent No. 5,473,599 issued to Li and Cole on December 5, 1995. Both of these references are incorporated herein by reference in their entireties and for all purposes.

If Home Agent 8 participated in a HSRP (together with other routers connected to segment 12), its failure would allow those nodes currently plugged into sub-network 12 to maintain their communications with internet 4. However, its failure would leave Mobile Node 6 stranded on network segment 14. HSRP has no mechanism for handling internet communications via Mobile IP. This is partly due to the fact that in Mobile IP it is not enough to simply have a standby router ready to take over as active router. The Home Agent (active router) must carry-out higher level functions required by Mobile IP such as keeping track of the locations (and associated Foreign Agents) of the various Mobile Nodes for which it is responsible. Similarly, a Foreign Agent must keep track of visiting Mobile Nodes and their associated Home Agents.

One redundancy mechanism for Mobile IP has been proposed. It goes by the acronym HARP which stands for Home Agent Redundancy Protocol. It was presented in an Internet Engineering Task Force memo of Chambless and Binkley entitled "Home Agent Redundancy Protocol" and having a URL of <ftp://ietf.org/internet-drafts/draft-chambless-mobileip-harp-00.txt>. This protocol provides for a redundant or "peer" Home Agent which is intended to contain a record of the Mobile Node locations stored in the primary Home Agent. While this proposed protocol does provide a redundancy mechanism for Mobile IP Home Agents, it has certain short comings. Notably, it does not make use of a widely installed redundancy protocol. Therefore, to implement HARP, many enterprises must undertake a rather significant change to its existing network solution. More importantly, HARP is concerned only with Mobile IP. Many other network functions such as Network Address Translation ("NAT"), IP security, Reflexive Access List, etc. all could profit from router redundancy. However, each of these has its own specific high level requirements (analogous to the mobility binding table required for Mobile IP). HARP cannot be easily extended to these Non-Mobile IP network

functions.

Unlike HARP, HSRP is a widely installed redundancy protocol. However, even if HSRP were implemented as the underlying redundancy protocol to implement Home Agent redundancy in a Mobile IP environment, there are several disadvantages associated with HSRP. First, HSRP enables only one router for a given network segment to act as the standby router at a single point in time. In other words, a redundancy scheme implemented using HSRP could include only a single backup Home Agent. Second, HSRP must be implemented on the same Local Area Network (LAN). It would therefore be beneficial if Home Agent redundancy could be implemented to enable multiple Home Agents to act as a backup for an active router. Moreover, it would be desirable if such a scheme could be implemented in a Wide Area Network (WAN) as well as a LAN. In addition, it would be beneficial if such a Home Agent redundancy scheme could be implemented regardless of the underlying routing redundancy protocol used to support such a scheme.

For the above reasons, an improved Home Agent redundancy protocol is required for Mobile IP.

SUMMARY OF THE INVENTION

Methods and apparatus for maintaining Mobile IP operation in a mobility

agent such as a Home Agent are disclosed. More particularly, this is accomplished through enabling information maintained by each mobility agent to be transferred among one another such that the information maintained by each is consistent. In this manner, multiple mobility agents may supply redundancy within a network.

5 In accordance with one aspect of the invention, a Home Agent (e.g., active Home Agent) sends a multicast message updating members of a virtual router group of a Mobile IP registration (or de-registration). First, a Mobile Node is registered and a registration entry is created in a mobility binding table for the Mobile Node. A
10 multicast message is then sent to a virtual router group to which the active Home Agent belongs and with which the active Home Agent shares a virtual IP address. The multicast message notifies the virtual router group of the registration. A similar process is performed when a Mobile Node is de-registered.

In accordance with another aspect of the invention, when a Home Agent participating in the redundancy group is initialized, it sends a multicast mobility
15 binding table request to the redundancy group indicating that bindings are requested. For instance the Home Agent may be active or non-active (e.g., standby, listening, or learning). The Home Agent may then receive bindings in response to the request and update its mobility binding table with the received bindings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A is a diagram of a Mobile IP network segment and associated environment.

FIG. 1B is a diagram of the Mobile IP environment of Figure 1A, depicting its limitations resulting from failure of a Home Agent.

FIG. 2 is a diagram illustrating an exemplary mobility binding table maintained by a Home Agent that may be updated in accordance with various embodiments of the invention.

FIG. 3 is a process flow diagram illustrating a method of implementing Home Agent redundancy in accordance with one embodiment of the invention.

FIG. 4 is a process flow diagram illustrating a method of initializing a Home Agent as shown at block 302 of FIG. 3.

FIG. 5 is a process flow diagram illustrating a method of downloading a mobility binding table as shown at block 304 of FIG. 3.

FIG. 6 is a process flow diagram illustrating a method of processing a registration request as shown at block 306 of FIG. 3.

FIG. 7 is a process flow diagram illustrating a method of processing binding table updates sent at block 610 of FIG. 6.

FIG. 8 is a process flow diagram illustrating a method of performing de-registration in accordance with one embodiment of the invention.

FIG. 9 is a process flow diagram illustrating a method of receiving and processing a de-registration update sent at block 804 of FIG. 8.

FIG. 10 is a diagram illustrating an exemplary mobility binding table update that may be sent by an active Home Agent to other Home Agents in its virtual router group upon receipt of a valid registration request from a mobile node.

FIG. 11 is a diagram illustrating an exemplary acknowledgement message that may be sent by a peer Home Agent to the active Home Agent in response to receipt of the mobility binding table update such as that illustrated in FIG. 10.

FIG. 12 is a diagram illustrating an exemplary mobility binding table request message that may be sent by a Home Agent.

FIG. 13 is a diagram illustrating an exemplary binding table update message that may be sent in response to a binding table request message such as that illustrated in FIG. 12.

FIG. 14 is a diagram illustrating an exemplary acknowledgement message that may be sent in response to the binding table update message such as that illustrated in FIG. 13.

FIG. 15 is a diagram illustrating an exemplary network device in which embodiments of the invention may be implemented.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be obvious, however, to one skilled in the art, that the present invention may be practiced without
5 some or all of these specific details. In other instances, well known process steps have not been described in detail in order not to unnecessarily obscure the present invention.

The present invention supports a mechanism for implementing Home Agent redundancy. This is accomplished by enabling a mobility binding table maintained
10 by a plurality of Home Agents to be maintained such that they are consistent with one another. Since all Home Agents within a particular redundancy group receive the same bindings, any one of these Home Agents may operate in place of a non-functioning Home Agent. In this manner, the ability of members of a redundancy group (e.g., HSRP group) that are in a non-active state such as a standby or “passive”
15 state (e.g., listen or learn) may be leveraged. This is particularly important since previously implemented schemes have not leveraged those Home Agents in a passive state such as the listen or learn state. While the following description is illustrated with respect to HSRP, such a scheme may be applied to members of a redundancy group using any redundancy protocol, where each member is either a primary Home
20 Agent (e.g., active) or a backup Home Agent (e.g., passive).

In the following description, the term “active” will be used interchangeably with the term “primary”, and the term “backup” will be used interchangeably with the term “standby.” A standby or backup router may, at times, be in a “passive” state. More particularly, members that are in a passive state typically gather information

without interacting or sending information to other members in the redundancy group. For example, with respect to HSRP, the term active refers to an active state, while the term non-active refers to a standby or passive (e.g., listen or learn) state. However, these states are merely illustrative, and the terms active and passive may be

5 interpreted generally as applied with respect to other protocols.

Each Home Agent within a particular redundancy group maintains its own internal mobility binding table. FIG. 2 is a diagram illustrating an exemplary mobility binding table that may be maintained by a Home Agent in accordance with various embodiments of the invention. As shown, each entry typically includes a

10 node ID 202 such as an IP address and a care-of address 204. In addition, each entry also preferably includes a granted lifetime and remaining lifetime.

There are various ways in which a redundancy scheme may be implemented. FIG. 3 is a process flow diagram illustrating a method of implementing Home Agent redundancy in accordance with one embodiment of the invention. Each Home Agent

15 in a particular redundancy group (e.g., virtual router group) is initialized at block 302 to enable each Home Agent to update its mobility binding table to contain bindings from other Home Agents that are members of the same virtual router group. One method of initializing the operation of a Home Agent is through sending a multicast binding table request and receiving mobility binding table updates in response to this

20 request from the members of the virtual router group. Such an initialization process will be described in further detail below with reference to FIG. 4.

Once a Home Agent is initialized, it is capable of responding to Home Agents requesting mobility binding table updates. More particularly, the Home Agent may respond to such mobility binding table requests regardless of whether the Home

Agent is active. For example, the Home Agent may be in a standby or listen mode. Thus, as shown at block 304, mobility binding table updates are sent in response to a received mobility binding table request. In addition, the Home Agent may update its list of Home Agents in its associated virtual router group using information provided in the mobility binding table request. One method of downloading a mobility binding table to a requesting Home Agent will be described in further detail below with reference to FIG. 5.

When a registration request is received by a Home Agent (e.g., Home Agent in an active or passive state) at block 306, the Home Agent sends a mobility binding table update to routers in the redundancy group to notify the routers of the new binding. One method of completing registration of a mobile node will be described in further detail below with reference to FIG. 6 and FIG. 7. Similarly, when a mobile node returns to its home subnet, de-registration is performed at block 308. Similarly, when the mobile node roams to a new Foreign Agent, de-registration may also be performed (in addition to sending a new registration request). One method of performing de-registration of a mobile node will be described in further detail below with reference to FIG. 8 and FIG. 9.

As described above with reference to block 302 of FIG. 3, a Home Agent may be initialized such that its mobility binding table contains bindings maintained by other Home Agents in its redundancy group. FIG. 4 is a process flow diagram illustrating a method of initializing a Home Agent as shown at block 302 of FIG. 3 in accordance with one embodiment of the invention. As shown at block 402, the Home Agent boots up and determines whether further initialization is necessary. More particularly, the Home Agent initializes the state of the mobility binding table = NOT

SYNCED at block 404. It then determines whether one or more Home Agents other than the Home Agent being initialized are in the virtual router group at block 406. If there are no other Home Agents in the virtual router group, the state of the mobility binding table = SYNCED at block 408. Otherwise, if it is determined that there are other Home Agents in the virtual router group, a multicast mobility binding table request is sent at block 410 to the virtual router group. More particularly, the multicast mobility binding table request indicates that bindings present in a mobility binding table maintained by one or more members of the virtual router group are requested. An exemplary mobility binding table request will be described in further detail below with reference to FIG. 12.

When each Home Agent in the virtual router group receives the multicast mobility binding table request, it updates its list of Home Agents in the redundancy group, since the multicast mobility binding table request identifies the Home Agent sending the mobility binding table request. As described above, the Home Agent may, for example, be active, or may merely be a standby or listening Home Agent. When a receiving Home Agent receives the multicast mobility binding table request, it sends one or more reply messages (e.g., unicast message) including bindings associated with one or more mobile nodes. It is important to note that all Home Agents in the redundancy group have an identical mobility binding table. Thus, rather than requiring each Home Agent in the redundancy group to reply to a multicast mobility binding table request, a designated Home Agent in the redundancy group may reply to the request. Such reply messages are received by the requesting Home Agent at block 412, which and are then used to update its mobility binding table. An exemplary binding table reply message will be described in further detail

below with reference to FIG. 13. This process continues at block 414 for the remaining bindings and reply messages. Once all bindings have been obtained from one or more Home Agents in its virtual router group (e.g., designated Home Agent), the mobility binding table state is changed to SYNCED at block 416. In this manner,

5 the Home Agent can track the state of its initialization and whether its mobility binding table is current.

Once a Home Agent has an updated mobility binding table, it may download bindings from its mobility binding table to other requesting Home Agents. FIG. 5 is a process flow diagram illustrating a method of downloading a mobility binding table as shown at block 304 of FIG. 3 in accordance with one embodiment of the invention.

10 As shown at block 502, the Home Agent receives a multicast mobility binding table request addressed to the virtual router group. As described above, the mobility binding table request includes a Home Agent identifier associated with the Home Agent that has sent the mobility binding table request. Thus, the Home Agent stores

15 the Home Agent identifier from the request in a list of Home Agents in the virtual router group at block 504. Through this list, the Home Agent may perform load balancing of registration requests it receives from various mobile nodes. Later, if the Home Agent sends a binding table update message, it may determine from this list from which Home Agents to expect an acknowledgement message.

20 Although a Home Agent need not be active in order to respond to a mobility binding table request, it should preferably have completed its initialization of its mobility binding table with bindings from other Home Agents in its redundancy group. Thus, at block 506, it determines whether its mobility binding table is initialized (e.g., synced). If its mobility binding table is not initialized, the Home

Agent sends a binding table reply with an error code at block 508. More particularly, the error code may indicate that the mobility binding table maintained by the Home Agent is not initialized with bindings from other Home Agents in the virtual router group.

5 If the mobility binding table of the Home Agent has been initialized, it sends one or more binding table reply messages at block 510. As described above, an exemplary binding table reply message will be described in further detail below with reference to FIG. 13. More particularly, these reply messages are preferably unicast messages addressed to the Home Agent in the virtual router group from which the

10 multicast mobility binding table request was received. As shown at block 512, one or more reply messages including one or more Home Agent identifiers identifying Home Agents in the virtual router group may be sent. For example, the Home Agent identifiers may be obtained from a list maintained by the Home Agent. More particularly, the Home Agent identifiers may be provided in a first reply message

15 followed by bindings provided in subsequent reply messages. Once the Home Agents in the virtual router group are identified, this enables the receiving Home Agent to perform load balancing of registration requests it subsequently receives from various mobile nodes. In other words, registration requests may be relayed to any of the identified Home Agents according to any load balancing algorithm. More

20 particularly, the receiving Home Agent may simply be in a listening state, but may perform load balancing when it subsequently becomes active. The receipt of such a list is important since underlying redundancy protocols do not provide periodic messaging to enable identification of other Home Agents to be performed. In addition, one or more reply messages including one or more bindings associated with

one or more mobile nodes obtained from the mobility binding table maintained by the Home Agent are sent at block 514. Reply messages are sent until all bindings are transmitted to the requesting Home Agent at block 516.

In addition to the sending of binding table reply messages, the Home Agent
5 also preferably verifies receipt of each binding table reply message by the intended recipient at block 518. This verification is preferably performed in parallel with the sending of the reply messages, as shown. In other words, it is not necessary to send all reply messages prior to verifying receipt of the reply messages. Thus, as shown at block 520, it is determined for each binding table reply message whether an
10 acknowledgement to the reply message has been received from the Home Agent. An exemplary acknowledgement message that may be sent in response to a binding table update message such as that illustrated in FIG. 13 will be described in further detail below with reference to FIG. 14. If an acknowledgement has not been received, the reply message is retransmitted to the Home Agent at block 522. Otherwise, when it is
15 determined that all reply messages have been received by the Home Agent, the process is completed at block 524.

A Home Agent in the redundancy group (e.g., active or passive) may receive and process a registration request received from a mobile node. FIG. 6 is a process flow diagram illustrating a method of processing a registration request as shown at
20 block 306 of FIG. 3. After the Home Agent receives a registration request at block 602, the registration request is processed according to RFC 2002 such that the mobile node is registered with the Home Agent at block 604. For instance, a registration entry for the mobile node is created in a mobility binding table. If the Home Agent determines at block 606 that there are no other Home Agents in the virtual router

group, the process ends at block 608. However, if there are other Home Agents in the virtual router group, a multicast binding table update is sent to a virtual router group to which the Home Agent belongs and with which the Home Agent shares a virtual IP address at block 610. The multicast binding table update notifies the virtual router group of the registration of the mobile node with the Home Agent. An exemplary binding table update message that may be sent will be described in further detail below with reference to FIG. 10.

Once the multicast binding table update has been sent, the Home Agent preferably verifies receipt of the update by all other Home Agents in the virtual router group. Thus, at block 612 the Home Agent determines whether an acknowledgement to the multicast message has been received from a Home Agent in the virtual router group. If an acknowledgement has not been received from a Home Agent (e.g., within a specified period of time), a message is retransmitted to the Home Agent notifying the Home Agent of the registration at block 614. More particularly, this message may be a unicast message since the other Home Agents in the virtual router group have already acknowledged receipt of the multicast message. Once an acknowledgement has been received from the other Home Agents in the virtual router group, the process ends at block 616.

As described above with reference to block 610 of FIG. 6, multicast binding table updates are sent to other Home Agents in the virtual router group. FIG. 7 is a process flow diagram illustrating a method of processing binding table updates sent at block 610 of FIG. 6 by a receiving Home Agent in accordance with one embodiment of the invention. A binding table update such as that described below with reference to FIG. 10 is received at block 702. More particularly, the binding table update

includes at least one binding associated with a mobile node that has registered with a Home Agent (e.g., active or passive Home Agent). The receiving Home Agent (e.g., active or passive Home Agent) then updates its mobility binding table to include the binding at block 704. An acknowledgement message indicating receipt of the binding table update is then sent at block 706. An exemplary acknowledgement message will be described in further detail below with reference to FIG. 11.

When a mobile node returns to its home subnet or has roamed to a new Foreign Agent, the Home Agent de-registers the mobile node (in association with the Foreign Agent from which the mobile node has roamed). Typically, when a mobile node roams from a first Foreign Agent to a second Foreign Agent, de-registration is performed with respect to the first Foreign Agent while registration is performed such that the mobile node is associated with the second Foreign Agent. Of course, there is a period of time during which a mobile node may receive a message after it has roamed from the first Foreign Agent but before it has completed de-registration with respect to the first Foreign Agent. However, the Home Agent can detect when its information is not current, since the mobile node cannot be simultaneously located at two different Foreign Agents. Thus, in order for the Home Agent to detect which registration is the current registration, it is possible to simply determine the order of the registrations of the mobile node.

However, in accordance with the present invention, since bindings are sent to other Home Agents as well as the primary Home Agent (e.g., active Home Agent), there is the added possibility outdated bindings may be propagated to other Home Agents in a redundancy group. FIG. 8 is a process flow diagram illustrating a method of performing de-registration in accordance with one embodiment of the invention

that alleviates the potential of propagating outdated bindings. Once the mobile node is de-registered with the Home Agent, the registration entry in the mobility binding table associated with the mobile node is updated (e.g., cleared) at block 802. A multicast binding table update is then sent to the virtual router group to which the Home Agent belongs and with which the Home Agent shares a virtual IP address at block 804. An exemplary multicast binding table update that may be sent will be described in further detail below with reference to FIG. 10. In this manner, the Home Agent notifies the virtual router group of the de-registration.

When a Home Agent (e.g., active or passive) receives a binding table update notifying it of a de-registration, it processes this information accordingly. FIG. 9 is a process flow diagram illustrating a method of receiving and processing a de-registration update sent at block 804 of FIG. 8 in accordance with one embodiment of the invention. When a Home Agent receives a de-registration binding table update indicating that a mobile node is no longer registered with another Home Agent at block 902, the receiving Home Agent updates its mobility binding table accordingly. More particularly, the receiving Home Agent determines whether a binding exists in its mobility binding table at block 904. If a binding exists as shown at block 906, the receiving Home Agent determines whether the de-registration request is more current than the binding. If the de-registration request is more current than the binding, the Home Agent updates its mobility binding table to remove the binding at block 908. In addition, the Home Agent removes tunnels and routes required to forward packets to the mobile node. If a binding does not exist in its mobility binding table, an entry is created in the mobility binding table for the mobile node at block 910 and the entry is marked as "deleted" at block 912. Then, if another binding table update for the

mobile node is subsequently received, the Home Agent can check the time stamp of the binding table update against the time stamp associated with the “deleted” entry to determine which entry is most recent. If the “deleted” entry is most recent, the Home Agent clears its binding table of the “deleted” entry. However, if the incoming

5 binding table update is more recent, the Home Agent updates its bindings to remove the “deleted” tag. This entry may be purged after a specified “safe” period of time at block 914 to enable memory to be freed for unnecessary entries.

As described above with reference to FIG. 6, when a registration request is received and processed by a Home Agent (e.g., active or passive Home Agent), the

10 Home Agent sends a mobility binding table update at block 610 to provide other routers in its virtual router group with this new binding. FIG. 10 is a diagram illustrating an exemplary mobility binding table update that may be sent by a Home Agent to other Home Agents in its virtual router group upon receipt of a registration request from a mobile node. This message contains information from the registration

15 request and in addition, two lifetime fields. In addition, type 1002 indicates that this message is a Home Agent redundancy message (e.g., binding information sync message). The two lifetime fields include a granted lifetime 1004 indicating a lifetime granted to the mobile node by the Home Agent and a remaining lifetime 1006 indicating a lifetime remaining until the mobile node’s mobility binding expires.

20 Remaining fields include information obtained from the registration request, such as a Home Address 1008 of the mobile node, Home Agent 1010 to which the registration request was transmitted, care-of address 1012 of the Foreign Agent to which the mobile node has roamed, and Identification 1014 for use in identifying the mobility binding table update message. The Identification field 1014 may be used to verify

and match valid acknowledgement messages, which will be described in further detail below with reference to FIG. 7.

As described above with reference to FIG. 7, when a mobility binding table update is received by a Home Agent (e.g., active or passive Home Agent), the receiving Home Agent preferably acknowledges receipt of this update as described above with reference to block 706. FIG. 11 is a diagram illustrating an exemplary acknowledgement message that may be sent by a peer Home Agent to the sending Home Agent in response to receipt of the mobility binding table update such as that illustrated in FIG. 10. Type 1102 indicates that this message is a Home Agent redundancy message (e.g., binding information sync acknowledgement message). In addition, a code 1104 indicates the result of the mobility binding table update message sent in response to receipt of a registration request. More particularly, the code 1104 indicates whether the registration update message was accepted by the receiving Home Agent, an unspecified reason for non-acceptance of the registration, whether the registration is prohibited for an administrative reason (e.g., authentication failed) , or that there are insufficient resources (e.g., memory) to the registration. In addition, the acknowledgement message includes the Home Address 1106 of the mobile node obtained from the received mobility binding table update message (e.g., bind information sync message) and the Identification 1108 obtained from the mobility binding table update message. The identification field enables the acknowledgement messages to be matched with the mobility binding table update messages and further protects against replay attacks of the messages.

As described above with reference to block 410 of FIG. 4, during initialization of a Home Agent, it may request bindings from other Home Agents in its virtual

router group. FIG. 12 is a diagram illustrating an exemplary mobility binding table request that may be sent by a Home Agent. This mobility binding table request is preferably a multicast message, but may also be a unicast message. Moreover, such a message is preferably sent by a peer Home Agent either prior to assuming a standby

5 role or after assuming the standby role. The type 1202 indicates that this message is a Home Agent redundancy message (e.g., binding table request message). In addition, a NumHA field 1204 may be used to indicate a number of Home Agent addresses for which the Home Agent is requesting that bindings be downloaded. Alternatively, the default may be to request that all bindings be downloaded, without specifying the

10 number of Home Agent addresses. An Identification field 1206 is used to identify the request. Moreover, Home Agent Address(es) 1208 may be used to indicate Home Agents for addresses that should be downloaded corresponding to the NumHA field 1204. More particularly, the Home Agent Address(es) 1208 may simply specify that all address bindings should be downloaded, in which case the NumHA field 1204

15 may be set to a pre-defined constant.

When a binding table request is received by a Home Agent (e.g., active or passive), a binding table reply message is sent as described above with reference to block 412 of FIG. 4 and FIG. 5. FIG. 13 is a diagram illustrating an exemplary binding table update message that may be sent in response to a binding table request

20 message such as that illustrated in FIG. 12. The Type 1302 indicates that this message is a Home Agent redundancy message (e.g., binding table sync message). In addition, Code 1304 may include a value indicating the result of receipt of the binding table request. For instance, the Code 1304 may indicate whether the binding table request message was accepted by the sending Home Agent, an unspecified reason for

non-acceptance of the message, whether processing of the message is prohibited for an administrative reason (e.g., mobility binding table has not been initialized) or that there are insufficient resources (e.g., memory) to process and/or adequately respond to the message. For example, as described above with reference to block 508 of FIG.

5 5, a reply message with an error code may be returned when the mobility binding table is not “SYNCED.” In addition, multiple binding table reply messages may be sent as described above with reference to blocks 510-516 of FIG. 5. Thus, NumBind field 1306 may be used to indicate the number of mobile node bindings included in a particular binding table sync message. In addition, a MoreBind field 1308 may be
10 used to indicate the number of mobile node bindings transmitted thus far (including this message). An Identification field 1310 may be used to identify this message so that it may be matched with a corresponding acknowledgement message, described below with reference to FIG. 14, as well as the corresponding binding table request message. One or more extensions 1312 may also be appended for transmitting each
15 binding.

As described above with reference to blocks 518-522 of FIG. 5, a Home Agent transmitting bindings to another Home Agent preferably verifies whether the receiving Home Agent has acknowledged those bindings. FIG. 14 is a diagram illustrating an exemplary acknowledgement message that may be sent in response to
20 the binding table update message such as that illustrated in FIG. 13. The Type 1402 indicates that this message is a Home Agent redundancy message (e.g., binding table sync acknowledgement message). In addition, Code 1404 may be a value indicating the result of receipt of the binding table sync message. For instance, the Code 1404 may indicate whether the binding table sync message and associated bindings were

accepted by the receiving Home Agent, an unspecified reason for non-acceptance of the message, whether processing of the message is prohibited for an administrative reason (e.g., authentication failure) or that there are insufficient resources (e.g., memory) to process and/or adequately respond to the message. The message may also include MoreBind field 1406 from the binding table sync message and the Identification field 1408 from the binding table sync message. Thus, the MoreBind field 1406 and the Identification field 1408 may be used, separately or in combination, to match the binding table sync acknowledgement message of FIG. 14 with the binding table sync message of FIG. 13. The Identification field 1408 may also be used to match the binding table sync acknowledgement message of FIG. 14 with the original binding table request of FIG. 12.

The above-described messages are merely illustrative. Thus, these messages may include additional or alternative fields. In addition, various features are not shown to simplify the illustrations. For instance, each Home Agent redundancy message sent by a Home Agent, such as those described above with reference to FIG. 10 through FIG. 14, preferably include a Home Agent-Home Agent Authentication Extension to be used for authenticating messages sent between the Home Agents. Thus, in order to provide such an authentication extension, the Home Agents preferably share a single key.

The present invention enables multiple Home Agents to provide redundancy, where each of the Home Agents preferably implements an underlying routing redundancy protocol. For instance, the routing redundancy protocol may be HSRP or Virtual Router Redundancy Protocol (VRRP).

Other Embodiments

Generally, the techniques of the present invention may be implemented on software and/or hardware. For example, they can be implemented in an operating system kernel, in a separate user process, in a library package bound into network applications, on a specially constructed machine, or on a network interface card. In a specific embodiment of this invention, the technique of the present invention is implemented in software such as an operating system or in an application running on an operating system.

A software or software/hardware hybrid implementation of the techniques of this invention may be implemented on a general-purpose programmable machine selectively activated or reconfigured by a computer program stored in memory. Such a programmable machine may be a network device designed to handle network traffic, such as, for example, a router or a switch. Such network devices may have multiple network interfaces including frame relay and ISDN interfaces, for example. Specific examples of such network devices include routers and switches. For example, the Home Agents of this invention may be implemented in specially configured routers or servers such as specially configured router models 1600, 2500, 2600, 3600, 4500, 4700, 7200, 7500, and 12000 available from Cisco Systems, Inc. of San Jose, California. A general architecture for some of these machines will appear from the description given below. In an alternative embodiment, the techniques of this invention may be implemented on a general-purpose network host machine such as a personal computer or workstation. Further, the invention may be at least partially implemented on a card (e.g., an interface card) for a network device or a general-purpose computing device.

Referring now to FIG. 15, a network device 1560 suitable for implementing

the techniques of the present invention includes a master central processing unit (CPU) 1562, interfaces 1568, and a bus 1567 (e.g., a PCI bus). When acting under the control of appropriate software or firmware, the CPU 1562 may be responsible for implementing specific functions associated with the functions of a desired network device. For example, when configured as an intermediate router, the CPU 1562 may be responsible for analyzing packets, encapsulating packets, and forwarding packets for transmission to a set-top box. The CPU 1562 preferably accomplishes all these functions under the control of software including an operating system (e.g. Windows NT), and any appropriate applications software.

CPU 1562 may include one or more processors 1563 such as a processor from the Motorola family of microprocessors or the MIPS family of microprocessors. In an alternative embodiment, processor 1563 is specially designed hardware for controlling the operations of network device 1560. In a specific embodiment, a memory 1561 (such as non-volatile RAM and/or ROM) also forms part of CPU 1562. However, there are many different ways in which memory could be coupled to the system. Memory block 1561 may be used for a variety of purposes such as, for example, caching and/or storing data, programming instructions, etc.

The interfaces 1568 are typically provided as interface cards (sometimes referred to as "line cards"). Generally, they control the sending and receiving of data packets over the network and sometimes support other peripherals used with the network device 1560. Among the interfaces that may be provided are Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces, token ring interfaces, and the like. In addition, various very high-speed interfaces may be provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, HSSI interfaces, POS interfaces, FDDI interfaces, ASI interfaces, DHEI interfaces

and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM. The independent processors may control such communications intensive tasks as packet switching, media control and management. By providing separate processors for the communications intensive tasks, these interfaces allow the master microprocessor 1562 to efficiently perform routing computations, network diagnostics, security functions, etc.

Although the system shown in FIG. 15 illustrates one specific network device of the present invention, it is by no means the only network device architecture on which the present invention can be implemented. For example, an architecture having a single processor that handles communications as well as routing computations, etc. is often used. Further, other types of interfaces and media could also be used with the network device.

Regardless of network device's configuration, it may employ one or more memories or memory modules (such as, for example, memory block 1565) configured to store data, program instructions for the general-purpose network operations and/or other information relating to the functionality of the techniques described herein. The program instructions may control the operation of an operating system and/or one or more applications, for example.

Because such information and program instructions may be employed to implement the systems/methods described herein, the present invention relates to machine readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable media include, but are not limited to, magnetic media such as hard disks, floppy

disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). The invention may also be embodied in a carrier wave travelling over an appropriate medium such as airwaves, optical lines, electric lines, etc. Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher level code that may be executed by the computer using an interpreter.

Although illustrative embodiments and applications of this invention are shown and described herein, many variations and modifications are possible which remain within the concept, scope, and spirit of the invention, and these variations would become clear to those of ordinary skill in the art after perusal of this application. For instance, the present invention is described as being implemented to enable Home Agents to operate to receive updates, regardless of their state. Although the terms passive and active are used in various exemplary descriptions above, the above-described embodiments may also be implemented with Home Agents in other states (e.g., standby). However, it should be understood that the invention is not limited to such implementations, but instead would equally apply regardless of the context and system in which it is implemented. Thus, broadly speaking, the operations described above may be used to enable redundancy with respect to other mobility agents, such as Foreign Agents. In addition, the above-described invention may be stored on a disk drive, a hard drive, a floppy disk, a server computer, or a remotely networked computer. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to

